CS 6840 – Algorithmic Game Theory (3 pages) 2/3/2012

Lecture 6: Utility Games

Instructor: Eva Tardos *Scribe: Jane Park (jp624)*

1 Announcements

- Scribe duty: try to get it done within a week while it's fresh.
- Final Projects
    1. Two types
        (a) Choose a favorite subarea from what we cover, "realize what we are missing" and "further the literature."
        (b) Incorporate game theoretic thinking into something else you're working on.
    2. Length: absolute max 10 pages. Min 5 or 6 pages.
    3. Partners: Try to work in pairs for the final project. Triples for the project is okay, if you can make it work.

2 Review: Price of Anarchy Bounds for Congestion Games

We derived Price of Anarchy bounds from the (λ, μ)-smooth inequality. A lot of you didn't buy this proof, so I'm doing two things to convince you: 1) We assume things we need. I'll provide examples where these things are actually true. 2) convince you that the bounds are sharp (we can't do any better).

3 Utility Games

Today we switch to utility games, another example of a game where the smoothness inequality results in games that have *Price of Anarchy* bounds. So far we've covered cost-minimizing games, but in other games like utility games, you derive benefits from participating in the game.

Definition. Utility game

- Players $1, \ldots, n$; player i has strategy set S_i
- Strategies give vector $s = s_1, \ldots, s_n$
- Player i's **utility** $u_i(s) \geq 0$: depends on strategy vector (what you're doing and what everyone else is doing)
- **Goal of player**: maximize utility
- s is Nash if $\forall i, s_i' \in S_i$:

$$u_i(s_i', s_{-i}) \leq u_i(s)$$

- A game is $(\lambda - \mu)$-smooth if:

$$\sum u_i(s_i^*, s_{-i}) \geq \lambda \sum_i u_i(s^*) - \mu \sum_i \mu_i(s)$$

Each $\sum$-term is utility (negative cost). We would like to place an upperbound on the cost, or lower bound on utilit at Nash. Intuition: if the optimal solution is better than the current solution, we hope someone discovers it via his/her $u_i(s_i^*, s_{-i})$ utility being high.

Note: for now we always assume that everyone knows everything, that is, we consider "full information games".

Theorem 1. If s is Nash equilibrium and s^* maximizes sum (s is at least as good as optimal):

$$\sum u_i(s) \geq \frac{\lambda}{\mu + 1} \sum_i u_i(s^*)$$

Proof.

$$\sum_i u_i(s) \geq \sum_i u_i(s_i^*, s_{-i}) \geq \lambda \sum_i u_i(s^*) - \mu \sum_i u_i(s)$$

$$(1 + \mu) \sum_i u_i(s) \geq \lambda \sum_i u_i(s^*)$$

□

Remark. This is a very different kind of game–no congestion/congestible elements.

Example. *Location Game*

Clients desire some sort of service, and k service **providers** position themselves to provide as much of this service as possible. Service providers(i) offer different **prices** p_{ij} to different **clients**(j). There is a service **cost** $c_{ij} > 0$ (not fixed, determined by location). Each client has associated **value** Π_i.

Client j selects the min price p_{ij}, and only if $\Pi_j \geq p_{ij}$. A client's benefit is $\Pi_j - p_{ij}$, and a client reacts to price only. We allow customers to make harmless changes (0-benefit) for mathematical simplicity.

Service provider i has customers A_i and benefit $\sum_{j \in A_i} p_{ij} - c_{ij}$. Provider locations are set, but prices can change often. Providers undercut each other to a point, then stop at equilibrium.

The natural outcome for given locations is that clients choose $\min_i c_{ij}$ (the nearest location) and the price offered is:

$$p_{ij} = \max(c_{ij}, \min(\Pi_j, \min_{k \neq i'} c_{kj}))$$

Explanation: $\min_{k \neq i'} c_{kj}$: the second cheapest location (the provider that you're afraid will undercut you). If the value of the customer's lower, charge that, but never charge less than your cost to serve the customer.

Note: This must be thought of as a multi-stage game where locations are fixed first, then prices are chosen.

Technical assumption: Cost never exceeds benefits.

$$\Pi_j \geq c_{ij}, \forall i, j$$

This assumption allows setting a much simpler rule for prices:

$$p_{ij} = \begin{cases} min_{i \neq k} c_{kj} & \text{by cheapest provider } k \\ c_{kj} & \text{by everyone else} \end{cases}$$

This assumption is helpful, and also harmless. There is no loss of generality since if cost does exceed value, we replace c_{ij} by Π_j. If this edge ever gets used for service, the price will be $p_{ij} = c_{ij}$ as all other edges have $c_{kj} \leq \Pi_j$ by assumption. This means that no one has any benefit from the edge: the user's benefit is $\Pi_j - p_{ij} = 0$ and the edge contributes to the provider's benefit by $p_{ij} - c_{ij} = 0$. So while our proposed solution can have such an edge with changed cost, we can drop the edge from the solution without affecting the anyone's solution quality.

Theorem 2. This game is also a potential game. Service providers are players, and their change in benefit is exactly matched by the change in potential function. We claim that the potential function that works here is social welfare–the sum of everyone's "happiness", and set Φ = social welfare.

Proof. Social welfare is the sum of all client and user benefits. Note: money (prices) does *not* contribute to benefit/welfare, but makes the economy run. Money cancels out due to its negative contribution on the client side and the positive contribution on the provider side. i_j is the location serving j.

$$\Phi = \sum_{j\text{:client served}} \Pi_j - c_{i_j j}$$

Change in Φ if provider i stops participating: let A_i be the set of users served by i. Each of them has to switch to the second closest provider now, so service cost increases; $\Delta\Phi = \sum_{i \in A_i} -c_{ij} + min_{k \neq i} c_{kj}$. The second closest provider was setting the price for $j \in A_i$, so $\Delta\Phi = \sum_{i \in A_i} -c_{ij} + p_{ij}$, exactly the value i had in the game.

To evaluate the change when i changes location, think of a two step process for provider switching location: 1) Go home (lose benefit) 2) Come back (gain benefit).

Note that the social welfare is the potential function. Will come back to it. □